Automating Vehicle Operations in Next-Generation Spacecraft: Human Factors Issues

Robert S. McCann* and Jeffrey McCandless.†
NASA Ames Research Center, Moffett Field CA 94035

Bruce Hilty‡
NASA Johnson Space Center, Houston, TX

NASA is designing a new space transportation system to replace the aging shuttles, return humans to the moon, and enable human exploration of more remote destinations. One of the fundamental requirements driving the design of this new system is that, due to speed-of-light limitations, many time-critical mission operations will have to be performed onboard the vehicle without real-time assistance from the ground. To meet this requirement, many mission management activities will involve real-time collaborations between crewmembers and onboard automation. We describe several human factors challenges that must be overcome to enable effective onboard human-machine teaming, including deciding on an appropriate functional allocation between human and machine, and developing and validating user interfaces to coordinate human and machine activities. We illustrate these principles with a concept for mixed initiative fault management on a next generation spacecraft.

I. Introduction

NASA's vision for future space exploration calls for the current Space Transportation System (STS) to be retired by 2010. A new Exploration Transportation System (ETS) capable of supporting crewed missions, first to the International Space Station (ISS) and other Low-Earth Orbit (LEO) destinations, then to the Moon, Mars, and beyond, will replace the STS. Key elements of the ETS include a Crew Exploration Vehicle (CEV), a Lunar Surface Access Module, and a Launch Vehicle. The current schedule calls for the CEV to assume responsibility for crew transport to and from the International Space Station (ISS) shortly after the shuttles are retired. The full suite of ETS vehicles and supporting elements must be completed in time for the first lunar mission sometime between 2015 and 2020.

How should these ETS vehicles be operated? While the question is certainly straightforward, the answer is complicated by several considerations. Any discussion of spacecraft operations must begin with the fact that human spaceflight is one of the most hazardous and unforgiving activities humankind has ever attempted. Human-rated spacecraft consist of very complex and often highly interactive physical systems, including multiple propulsion systems, a flight management system, electrical and mechanical power generation and distribution systems, and environmental control and life support systems. Particularly during the dynamic phases of flight, such as launch/ascent and entry, these systems must perform to precise operational specifications often under extremely harsh physical conditions. Thus, virtually any concept for interacting with and operating these systems involves some risk to crew safety and mission success.

Fortunately, over the last 45 years NASA has flown 140 LEO and lunar missions. These missions have provided the mission operations community with abundant opportunities to develop and refine operational concepts, steadily reducing operational risk. By now, any fundamental change to these operations, even if the name of modernization

*Group Lead, Intelligent Spacecraft Interface Systems Laboratory, Human Factors Research and Technology Division, Mail Stop 262-4, Moffett Field CA 94035.
† Research Psychologist, Human Factors Research and Technology Development Division, Mail Stop 262-4, Moffett Field, CA 94035.
‡Deputy Chief, Advanced Operations and Development Division, Mail Code DV, NASA Johnson Space Center, Houston, TX 77058.

and improvement, has the potential to expose new operational vulnerabilities and actually increase risk. As a result, the operational community sets a very high bar for such changes.

Given the stakes associated with human spaceflight, this conservative approach is both prudent and necessary. However, human spaceflight is at a pivotal crossroads. ETS missions are eventually going to require serious, fundamental changes to vehicle operations. In order to understand why changes are required, and what form they will take, a brief overview of current shuttle operations is necessary.

II. Current Operations

There are two distinct classes of spacecraft operations, those associated with the dynamic flight phases, such as ascent and entry, and those associated with the more quiescent on-orbit phase. On-orbit operations revolve around a daily schedule of crew activities, such as payload deployment, routine systems maintenance, equipment checkout, science experiments, and the like. We will have more to say about these operations in Section IV. For now, we will focus on the more dynamic phases, where the most dangerous systems are operating and the risks to crew and mission are greatest. To reduce these risks as far as possible, much of the operational focus is on acquiring, processing and interpreting real-time data concerning vehicle flight parameters (e.g., attitude, acceleration, trajectory) and the health and functioning of the active systems. To obtain these data, each system (along with critical structural components, such as the wings) is heavily instrumented with sensors that continuously generate numeric readings of key operating parameters, such as temperatures, pressures, flow rates, and accelerations. The most critical parameters are usually instrumented with more than one sensor to protect against individual sensor failures. In fact, the total number of sensors (and therefore, the total number of data sources) is over 2000.

A. Onboard Operations

Explicit information processing requirements exist for these data. The most basic requirement is that the data has to be monitored. If any parameter moves outside of its normal range of values, this "out-of-limits" condition must be detected, the cause of the condition must be identified and, if the cause is determined to be a genuine systems malfunction, appropriate remedial actions must be taken. Figure 1 shows the sequence of fault management activities that accompany actual systems malfunction in the shuttle cockpit. The vehicle's Caution and Warning (C&W) System, described in more detail below, detects an off-nominal sensor reading, sounds a cockpit alarm, and generates a flashing fault message on a cockpit display. A crewmember first silences the alarm by pressing the master alarm button, and then stabilizes the fault message by pressing the "acknowledge" key on the cockpit keyboard. The crewmember then reads the fault message, or several messages if the problem has "daughter" faults associated with it, and determines which fault message pertains to the root cause of the problem. Then, he or she locates the information pertaining to the malfunction in one of several paper flight data files, or on cue cards. To understand this information, generally a mix of troubleshooting activities and fault management procedures, the crewmember must decode what amounts to a cryptic pseudo-language consisting of specialized symbols, abbreviations, and both space- and form-based coding (i.e., line indentations, etc). The procedures themselves typically take the form of one or more switch throws that change the operating mode of the system in question, sometimes to further clarify and identify the source of the malfunction, sometimes to take advantage of built-in systems redundancies to recover nominal system function. The crewmember has to locate the appropriate switches from the hundreds of switches spread around the cockpit periphery, and then toggle them to the position indicated by the procedure. Finally, the crewmember must verify that the procedure has "safed" the system.

The human resources onboard the shuttles fall far short of what is needed to meet these data processing and fault management requirements. Only two crewmembers have visual access to the data on cockpit displays. It goes almost without saying that two "pairs of eyes" are insufficient to process thousands of individual data points in real time. Beyond that, only a small fraction of the sensed data can be viewed on cockpit displays. Of that fraction, display real estate limitations dictate that only an even smaller fraction can be viewed at any one time. If a crewmember wants to view all available data on a system, he or she has to navigate through several successive display formats. Processing of information on the individual display formats is often slow because the displays are poorly organized and highly cluttered, taking the form of closely-spaced tables of alphanumeric data that require considerable mental translation to infer the current operational status or functional mode of the system. And finally, during dynamic flight phases, any processing of vehicle systems data or work on an actual malfunction competes with urgent requirements to process other forms of data. A recent investigation of astronaut cockpit scanning patterns[1] revealed that, during a nominal ascent, astronauts spend between 25% and 50% of their time examining systems parameters on systems status displays; the majority of their time is spent examining displays of flight or mission-related information[2].

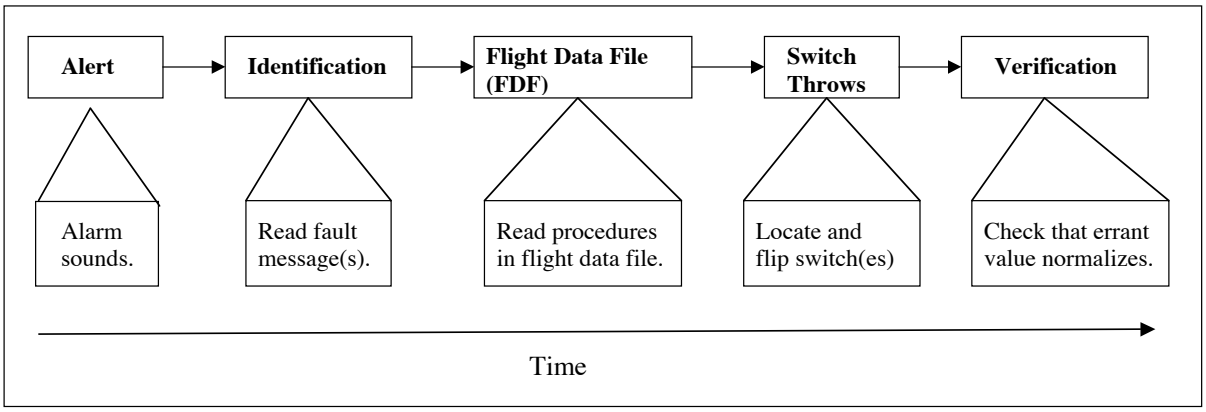

Figure 1. Fault Management Stages onboard the Shuttle

Since the onboard human information processing resources are insufficient to satisfy the data and information processing requirements of dynamic flight, what about the other onboard processing resource, the shuttle's five General Purpose Computers (GPCs)? Do they supply the needed additional resources? Unfortunately, the answer is no. The GPCs are of late 1970's vintage, with only rudimentary processing capabilities by today's standards. During dynamic flight, the computers are almost fully occupied with performing computations relating to vehicle guidance and flight control. Most of the small amount of capacity remaining is devoted to the C&W system, which serves as the primary communications link between the vehicle and the crew concerning systems health and functioning. The C&W system includes limit-sensing software that constantly compares selected sensor readings against preset upper and lower values. As we noted earlier, if a reading moves outside of either value, the out-of-limits condition is signaled in the cockpit by some combination of auditory and visual alarms. In addition, the alarm is typically accompanied by a fault message to help the crew identify the source of the problem.

By directing crewmembers' attention to an off-nominal reading, the C&W system performs a valuable function, automating much of the work involved in parameter monitoring and off-nominal detection that would otherwise completely overwhelm the capabilities of the crew. However, the system also has limitations that seriously restrict its usefulness. Since the limit-sensing software monitors each sensor individually, the system has no ability to discriminate a genuine off-nominal value from a value generated by a failed sensor (or by a failure in a signal processing component of the data processing system). More seriously, many failures, particularly in the Electrical Power System, generate off-nominal readings from a wide variety of subsystem components. This leads to a proliferation of C&W annunciations and fault messages that can seriously interfere with crew's attempts to build situation awareness and understand the source of the problem. The C&W system provides no further assistance with the activities involved in working a malfunction after it has been identified.

B. Ground Operations

With neither human nor machine-based resources adequate to cover the data and information processing requirements of the dynamic flight phases, how are these requirements actually met? The answer is via a telecommunications infrastructure that enables near real-time voice communications and data sharing (telemetry) between the vehicles and the ground. Much of the data generated by the onboard sensors is included in the telemetry stream and displayed to flight controllers and other subject matter experts at the mission control facilities. The telemetry stream enables ground-based subject matter experts to satisfy most of the requirements for data monitoring, detecting off-nominal parameter values, and making "root-cause" malfunction determinations that the vehicle cannot. Once a malfunction has been diagnosed, subsequent recovery operations take the form of tightly coupled collaborations between the ground and the crew.

Other time-and-safety-critical operations are equally dependent on real-time ground support. For example, each STS mission involves thousands of man-hours of advance planning to determine parameters such as the vehicle's flight performance envelope and expected trajectory. These computations are based on mission-specific variables such as vehicle payload and vehicle center-of-gravity. The problem is that these variables are all calculated prior to flight, and are subject to change in the event of an in-flight anomaly, such as an unexpected weather condition or structural breach. There is very little onboard capability to perform the computations necessary to support in-flight mission replanning in response to such anomalies, leaving the crew wholly dependent on the ground for such activities.

C. ETS Operations

Even the short description above makes clear how dependent the current operational concept is on near-real time telemetry and communications links between the ground and the vehicle. However, these links are available only for missions within the Earth-Moon system. On interplanetary missions, the vehicles will be so far away that speed-of-light limitations will effectively eliminate near real-time links with the ground. Consequently, ETS missions are eventually going to require a more *autonomous* concept of mission operations in which the most time-critical activities, particularly those associated with the dynamic flight phases, will have to be performed *on the vehicle itself*.

This operational change represents an enormous challenge. ETS vehicles will be faced with many of the same limits on onboard human processing resources that exist in today's vehicles. The crew complement is likely to be capped at six, still not nearly enough "pairs of eyes" to scan and digest all the sensor data in real time. ETS vehicles are going to resemble the capsules and modules of the Apollo era more than the Shuttles. Thus, if anything, the display real estate available in the cockpit will be less than on the Shuttles. To avoid overwhelming the crew, many of the requirements for data and information processing will have to be met by machines.

Are today's machines up to the challenge? There is good reason to believe that they are. In the time since the shuttles were developed, advances in computing and information technologies have been steadily chipping away at the original "driver" for ground-centered operations, lack of onboard mission management capability. Probably the single most important limitation, insufficient onboard computing capability, has been rendered moot by several iterations of Moore's Law, coupled with even more dramatic increases in the capacity of portable data storage devices. As for software development, sophisticated health-management systems exploit today's computing horsepower to process multiple data streams, and detect off-nominal data patterns, in real time. Model-based reasoning engines can associate off-nominal data patterns with the operational mode of the affected system, enabling automated fault diagnoses and "root-cause" failure determinations. Process controllers can automatically reconfigure the operational mode of a system (i.e., perform fault recovery procedures) and assess whether the reconfiguration has been successful. Beyond the health management arena, advanced flight management algorithms can respond to dynamic flight conditions and perform real-time flight replanning exercises, such as determining abort options in the event of a propulsion system malfunction during ascent. Planning and scheduling tools currently under development should soon be able to automatically generate schedules of crew activities, a nontrivial exercise that must take into account multiple simultaneous resource constraints. And, last but not least, human-centered user interface design principles and user interface technologies offer new capabilities to organize and present information to the crew in a way that maximizes human information processing capabilities and supports flexible forms of human-automation interaction.

1. Technology Infusion: Scheduling Issues.

Of course, the ETS will not be called upon to support an interplanetary mission for several decades. In the nearer term, CEV missions will be confined to LEO, ferrying crews and supplies to and from the International Space Station, and then will support crewed missions to the Moon. The existing, ground-centered operational concept would certainly suffice for these missions. Do we even have to worry about enhancing onboard mission management capabilities for them, particularly when the associated operational changes may expose the crews to risks that don't exist in current operations? We offer the following observations on this complex issue. There is no doubt that the CEV will incorporate some of today's advanced information technologies and supporting infrastructure; for example, it would be absurd to restrict the CEV's onboard computing capabilities to Apollo-era levels in a slavish attempt to adhere to Apollo-era operations. But any infusion of these technologies is an open invitation to enhance onboard mission management capabilities; indeed, what else will the technologies be used for? In addition, although, as we have noted, the existing "ground-centered" operational concept has the benefit of many years of experience in reducing operations-related risk, the concept is also expensive, inefficient, highly resource intensive, and so cumbersome that maintaining operational safety margins is a constant challenge. It is worth noting that, even though the Space Launch Initiative and Orbital Space Plane programs only targeted LEO operations, much of the conceptual and technical work performed for the programs involved leveraging modern information technologies to reduce the cost, improve the efficiency, and increase the safety margins of vehicle operations.

Other considerations also motivate an early move toward more autonomous operations. Even modest steps in this direction raise issues that directly impact the design requirements for a large number of vehicle components and systems. These include, but are not limited to, the architecture of the onboard data and information processing system, beginning with the placement and number of systems and structures sensors (decisions that are very difficult or impossible to modify via retrofitting); the nature and functioning of the vehicle caution and warning system, an issue we will take up below; user interfaces to support various human-machine function allocations (levels of automation); and crew interactions with mission control. Indeed, so large is the design space relating to operational issues that even the best and brightest designers are unlikely to produce optimized designs on their first iteration. An

early development phase is needed to rapidly prototype and evaluate various design concepts and their supporting infrastructures. Adding to the urgency of these activities, one of the most important "lessons learned" from automating operations in airline cockpits is that even the most carefully thought-out concepts for human-automation interaction contain unforeseen opportunities for nonstandard or unanticipated interactions that produce human error. An ongoing program of operational testing and evaluation is required to uncover these hidden "Gotchas", and modify the operational concept to eliminate them. And, last but not least, we return to the sheer difficulty and risk associated with developing and certifying operational concepts for crewed space vehicles. The nearer-term LEO and Lunar missions represent critical operational opportunities to prudently select technologies, mature them to the point where they can migrate onboard the vehicle, and thoroughly test and validate them in valid operational settings where contingency procedures (even the moon is only a couple of days journey from Earth) exist in case these new operations have a unexpected mission impact. By the time they are implemented on an interplanetary mission, where contingency procedures are few and far between, these concepts will have been thoroughly tested and validated.

D. Guidelines for Human-Automation Interaction

As we noted, to avoid overwhelming the crew, many of the requirements for data and information processing will have to be performed by onboard automation. Indeed, a legitimate issue is whether to try and use onboard automation to eliminate human involvement in real-time operations (including responses to emergencies) entirely, giving full responsibility to machines. There are several compelling reasons why this is not an appropriate operational target. First and foremost, entrusting the lives of a crew to software systems is inherently risky. Hardware or software failures are more common in space than on the ground, in part because space-based platforms are vulnerable to radiation-induced "single event upsets." Crewmembers are unlikely to trust software tools to the point where they cede all control over emergency operations. In addition, full automation is simply not the optimal way to utilize onboard human and machine resources. Crewmembers are (and should continue to be) trained in spacecraft operations and the architecture and functioning of vehicle systems until they are subject matter experts in their own right. Taking them out of the loop amounts to a decision to waste valuable onboard expertise. Equally important, humans and machines bring different capabilities and different vulnerabilities to bear on crucial real-time operations[3]. These capabilities are frequently complementary, with strengths in one compensating for weaknesses in the other. For example, humans are nondeterministic processors, which gives them a fluid reasoning capability that helps them solve problems in novel situations. Computing and related information technologies are still quite brittle in the face of the unexpected. They can, however, monitor, process, and recognize patterns in numeric data far faster and more accurately than people. In conjunction with today's advanced electronic display devices, computers can also organize and display information far faster and with much more flexibility than is the case with non-electronic information sources.

Given these considerations, we believe the key to enhancing onboard mission management capabilities on ETS vehicles is an operational concept in which crew and onboard intelligent systems (immobots[4]) partner with each other, working mission operations in an integrated, cooperative manner. In the remainder of this article, we consider some of the issues of human-machine interfaces and modes of functioning that must be addressed in order to achieve a workable concept for onboard human-machine teaming. Fortunately, many of the issues that need to be addressed to support effective human-machine "teaming" have been documented for quite some time[5]. These issues include, but are not limited to:

- Ensuring crew visibility into automated functioning. Automation is deemed "clumsy"[6] if the workings of the automation are opaque to the human. User interfaces must be designed that allow the crew to make determinations such as, "is the automation itself "healthy"? Is it performing in a manner consistent with what I know about how it functions, what kinds of computations it is performing, and what the outcomes of those computations are?

- Determining a functional allocation between humans and machines that A) capitalizes on the strengths and capabilities of both humans and machines, thereby optimizing the capabilities of the joint human-machine system, and B) avoids the "out-of-the loop unfamiliarity" (OOTLUF) problem[5,7]. The functional allocation needs to strike a balance between the reduction in workload that automation makes possible with the potential loss of situation awareness that can occur when the machine performs operations without sufficient human oversight and involvement. There are numerous examples from today's highly automated aircraft cockpits of serious consequences when crewmembers are insufficiently involved in an automated operation, and are suddenly called upon to deal with the consequences of that operation[7,8].

E. Backup Capability and Redundancy Requirements.

The uniquely hazardous nature of spaceflight magnifies the importance of avoiding these problems, and brings additional requirements for the design of human-machine systems. We have already noted the great complexity of spacecraft systems, and the harsh physical environments that they operate within during dynamic flight phases. Spacecraft systems are therefore much more vulnerable to mechanical failure than the systems onboard, say, a commercial aircraft. Indeed, much of the engineering complexity of the onboard systems is due to the need to build in operational redundancies so that if a component fails, a backup operational mode exists that restores or maintains full system functionality.

2. Hardware and Software Requirements.

The same stringent requirements for backup capability and redundancy extend to onboard automation. Several forms of redundancy are available, beginning with the hardware and software itself. For example, on the shuttle, guidance, navigation, and flight control functions are almost fully automated during ascent and entry. If something happens to either the flight software itself, or the hardware on which the flight software is running, the vehicle is immediately in great danger. Fully recognizing this vulnerability, the original shuttle designers had two separate contractors develop independent flight software systems. One company's software was designated primary, the other secondary. The primary software system, housed on four of the five GPC's, has nominal control over the vehicle. If the primary system were to fail, due to some combination of software or computer failure, the backup system, loaded on the fifth onboard GPC, can assume the most essential flight control functions.

There is an important human factors component to this redundancy. On ascent and entry, the two software systems continuously and independently compute critical flight parameters, such as vehicle attitude, which the backup system requires in case it has to engage. In addition, at two minutes into flight, additional guidance parameters are redundantly computed, such as the exact time to shut down the main engines to achieve the targeted orbital insertion point. These values, along with any real-time discrepancies between the two systems' computations of vehicle attitude, are continually displayed to the crew. By checking and crosschecking these values, the crew can continuously determine the health of each software system and the veracity of its navigation and guidance computations. As we move into a new operational environment, where onboard software becomes responsible for a much larger set of data and information processing activities, similar software and hardware redundancies should be built into these computations, and the crew will need similar "crosschecking" capabilities on their cockpit displays.

3. Crew as Backup

Redundancy requirements don't stop with the software and hardware. To protect against a general breakdown in the data processing system, the crew should have full capability to perform automated functions manually (ideally, the reverse would also be true; the automation should be able to act as "backup" in the event of a crewmember "malfunction". We shall return to this rather tricky issue when we discuss long-duration mission operations in Section IV). But this requirement exposes a human factors conundrum. In order to function effectively as a backup, a crewmember must retain the skill set needed to perform nominally automated functions. If a procedure or operation is always automated, crewmembers will never actually perform it, and will lack the skill set necessary to meet the backup requirement if the need arises. This issue can be dealt with through training scenarios that simulate automation failures that require a reversion to manual operations. However, another option is to design the crew-automation functional allocation and supporting interfaces in such a way that although the automation actually performs the operation, allowing the workload-reduction benefits of the automation to be realized, the human is kept "in the loop" in a manner that continuously reinforces the skill set necessary to perform the function manually.

F. A Concept for Crew-Automation Interaction: Real-Time Fault Management

We will now develop a concept for onboard human-automation interaction to illustrate how functional allocation determinations, and supporting user interface designs, can satisfy the various requirements for effective human-automation partnering identified in the previous section. We selected onboard fault management as our operational example for several reasons. As we have already seen, today's fault management operations (Figure 1) include several activities that place high demands on the crew's time and processing resources. With only limited forms of assistance available from the C&W system, crewmembers are highly reliant on real-time ground assistance when working most malfunctions. Thus, any effort to achieve more autonomous operations has to give fault management a high priority. Second, the state of the art in health management technology goes far beyond the limit-sensing and fault-messaging capabilities of the current C&W system. From an operations perspective, today's technology provides an opportunity to transform C&W into a "decision support" system capable of assisting the crew with every phase of the fault management process[9], greatly reducing the need for ground assistance.

Figure 2 provides a representative architecture for a "state of the art" health management system. The system assigns specific computational functions to a series of data-and-information processing "layers". Level 1, Signal Processing and Condition Monitoring, encompass algorithms that perform time-series analyses on continuously varying sensor data to detect off-nominal trends and discriminate nominal from off-nominal patterns. The next level, Health Assessment, encompasses rule-based and/or model-based reasoners that make root-cause diagnoses of off-nominal patterns. Finally, the Recovery and Safing level consists of a "smart" reactive planner that 1) determines what procedures are required in order to achieve the desired goal state (typically, a reconfiguration of the system's operating mode to restore nominal system function), 2) determines the correct sequence of procedures to achieve that state, 3) physically commands the procedures, and 4) via feedback from sensor data, determines whether the procedures have been carried out and whether the desired systems reconfiguration has been achieved.

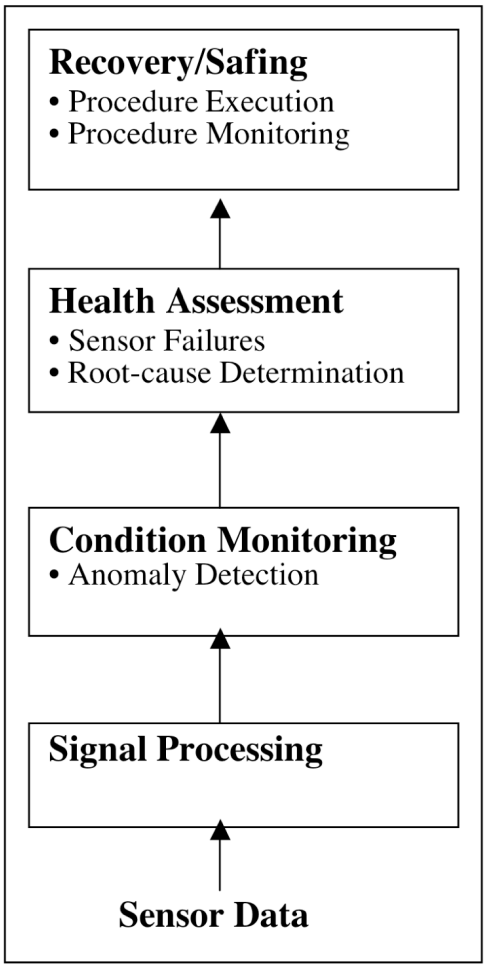

Figure 2. Systems Health Manager

The health management system depicted in Figure 2 has the ability to automate virtually all of the fault management activities illustrated in Figure 1 (current operations). Mindful of the dangers of over-automation, McCann and McCandless[10] argued that the optimal functional allocation between crew and health manager gives the automation responsibility for data monitoring, detecting off-nominal conditions, identifying the root cause, and executing the appropriate procedure(s). However, a crewmember maintains overall control of the process by having to "give permission" to carry out a procedure. The automation cannot perform the procedure until it receives this permission.

4. Making Automation Activities Transparent to the Crewmember

The user interfaces to support this teaming concept have several design requirements. One of the most important is to provide some indication to the crewmember of the health and functioning of each level in the fault management system. In the Shuttle cockpit, the equivalent of Level 1 is the limit-sensing software. When a sensor value exceeds its upper or lower limit for a set number of sequential readings, the out-of-limits value is signaled to crewmembers by changing the color of the normally white digital values (and adding an up or down arrow or other symbol beside the value) on the cockpit system summary displays. By contrast, Level 1 encompasses much more complex data processing algorithms, such as time-series-based analyses of the fluctuations (variance) in sensor readings over time. Such variance is a normal component of parameter data feeds, and analyzing that variance is a powerful tool for classifying nominal versus off-nominal modes of system functioning[11]. How might we capture the continuous operation of these algorithms on a system summary display without overloading the observer with too much information? One possibility is to render the digital values in grey, and code the variation in these values as subtle changes in brightness (brighter or dimmer levels of grey) as the values vary in real time. A to-be-determined bandwidth would be assigned to each displayed parameter. Deviations within the bandwidth would be coded by subtle changes in the brightness of the digits, slightly brighter for values above the mean, and slightly dimmer for values below. Again, these changes would only affect brightness along the grey scale; they would convey that normal signal variance was being measured, and also that Level 1 algorithms were "alive", processing the variance, and determining that the real-time values were falling within normal operational limits. If Level 1 determines that readings are beginning to deviate from nominal (outside of the normal bandwidth), all sensed values contributing to that assessment would change color. Yellow values would signal that an abnormal trend has been detected; red values would signal that an abnormal state has been confirmed. In this way, the basis for nominal and off-nominal color-and-intensity coding would change from today's system, based entirely on limit sensing, to reflect computations that integrate information across time and individual sensor feeds.

The next algorithmic layer is Health Assessment. Again, we would like a user interface that conveys some indication of the reasoning behind a root-cause determination. However, this issue is complicated by the plethora of algorithms and computations employed by different reasoners; some are strictly rule-based, others take a sensor fusion approach by subtracting out expected values and then matching the pattern of residuals against known failure modes. Since the user interface to make these operations transparent to the observer will depend on the specifics of the underlying computations, we will not attempt to provide a generic interface solution here.

The final level of the health management system is Procedure Retrieval and Execution. This represents one of the most significant extensions to cockpit automation capabilities over current capabilities, involves the tightest interactions between crewmembers and onboard automation, and (arguably) places the greatest demands on user interface design. The primary requirement for the interface is that it enables and coordinates the "permissions-based" concept for procedure execution. Arguably the most obvious candidate for this interface is a dedicated fault management display patterned after the systems and fault management displays developed for the Boeing 777 and Airbus "A" series of glass cockpit aircraft. When a malfunction is detected on an A320, for example, the appropriate set of procedures automatically appears in written form in a dedicated (normally blank) section of the ECAM display page. In addition, a "systems synoptic" appears on a separate Systems Page. The synoptic depicts components of the system affected by the malfunction in a spatial layout that matches the crews' mental model of system architecture and system functioning. The crew completes each procedure manually, and keeps track of the status of each procedure (completed versus not completed) via a checkmark that fills a box to the right of each procedure when completed. In addition, as the crew performs each procedure, the synoptic changes to display the new system configuration.

"Synoptic" systems displays are very helpful when determining which part of the system has failed and what a degraded mode of operation might be[12]. Recognizing these benefits, NASA Johnson Space Center recently completed a Cockpit Avionics Upgrade (CAU) project to reduce crew workload, enhance their situation awareness, and improve their performance. Among their other activities, CAU participants completely redesigned the shuttle's systems information displays, incorporating many of the synoptic features of the modern glass cockpit aircraft displays and developing luminance and color-based coding schemes to depict key aspects of system mode and system functioning. Our working assumption is that system summary displays on ETS vehicles will resemble the CAU redesigns much more than they will resemble the existing shuttle designs. Following Malin, et al.[13], McCann and McCandless[10] advocated a fault management display format that embeds procedural information right into the system schematic. This approach has the advantage that the dedicated fault management display can recapitulate essential elements of the system summary display, thus supporting rapid cross-referencing of information across the two formats.

Figure 3 depicts a hybrid design that incorporates elements of both a written electronic checklist and the "embedded-in-synoptic" approaches. The figure depicts a rather complex failure situation involving the helium supply system to one of the three engines that make up the shuttle's Main Propulsion System. When an engine is operating nominally, helium flows out of the tank at the top of the figure and then splits into two redundant legs, each with a separate pressure regulator, before rejoining and flowing to the engine in question (where the helium continuously pressurizes a seal in the engine's high pressure oxidizer turbopump). Following design criteria established by the CAU project, valves are depicted as circles with an embedded line. When the valve is open, the embedded line is flush with the rest of the line, and the entire line is bright white, indicating that helium is currently flowing. We also see that the Isolation Valve for the right-hand (B) Leg is open, and (as we would expect) helium is flowing through Leg B. However, the symbol for the Leg A Isolation Valve is colored red, the interior line is perpendicular to the flow, and the Leg A line below the valve is colored dark grey (signaling no flow). Together, these codes indicate that Isolation Valve A has failed to the "close" position. But the figure also indicates a second failure within the system. The "dP/dT" value in the upper right-hand corner is colored yellow, which indicates that helium is being depleted at a higher than nominal rate from the tank, that is, the valve failure has been compounded by a leak somewhere in the system. This set of circumstances conforms to a "Non-Isolatable Helium Leak" condition, indicated at the top of the display, with two associated procedures. When the helium tank pressure falls to a target value of less than 1150 PSI, a manifold connecting the engine's helium supply system with a backup helium supply system must be opened, to keep helium flowing to the engine for as long a period as possible. Then, to avoid any chance of full depletion of the helium supply, the engine must be shut down when the vehicle reaches an inertial velocity of 23,000 feet per second. We assume that the health management system is capable of estimating the time remaining before these conditions will be satisfied, and also that the tank pressure will deplete to the target value (1 min 12 sec from the present) before the vehicle reaches engine shut down velocity (almost 4 min from the present). Thus, the interconnect open procedure is first in the written procedures section (lower half of the display), and a

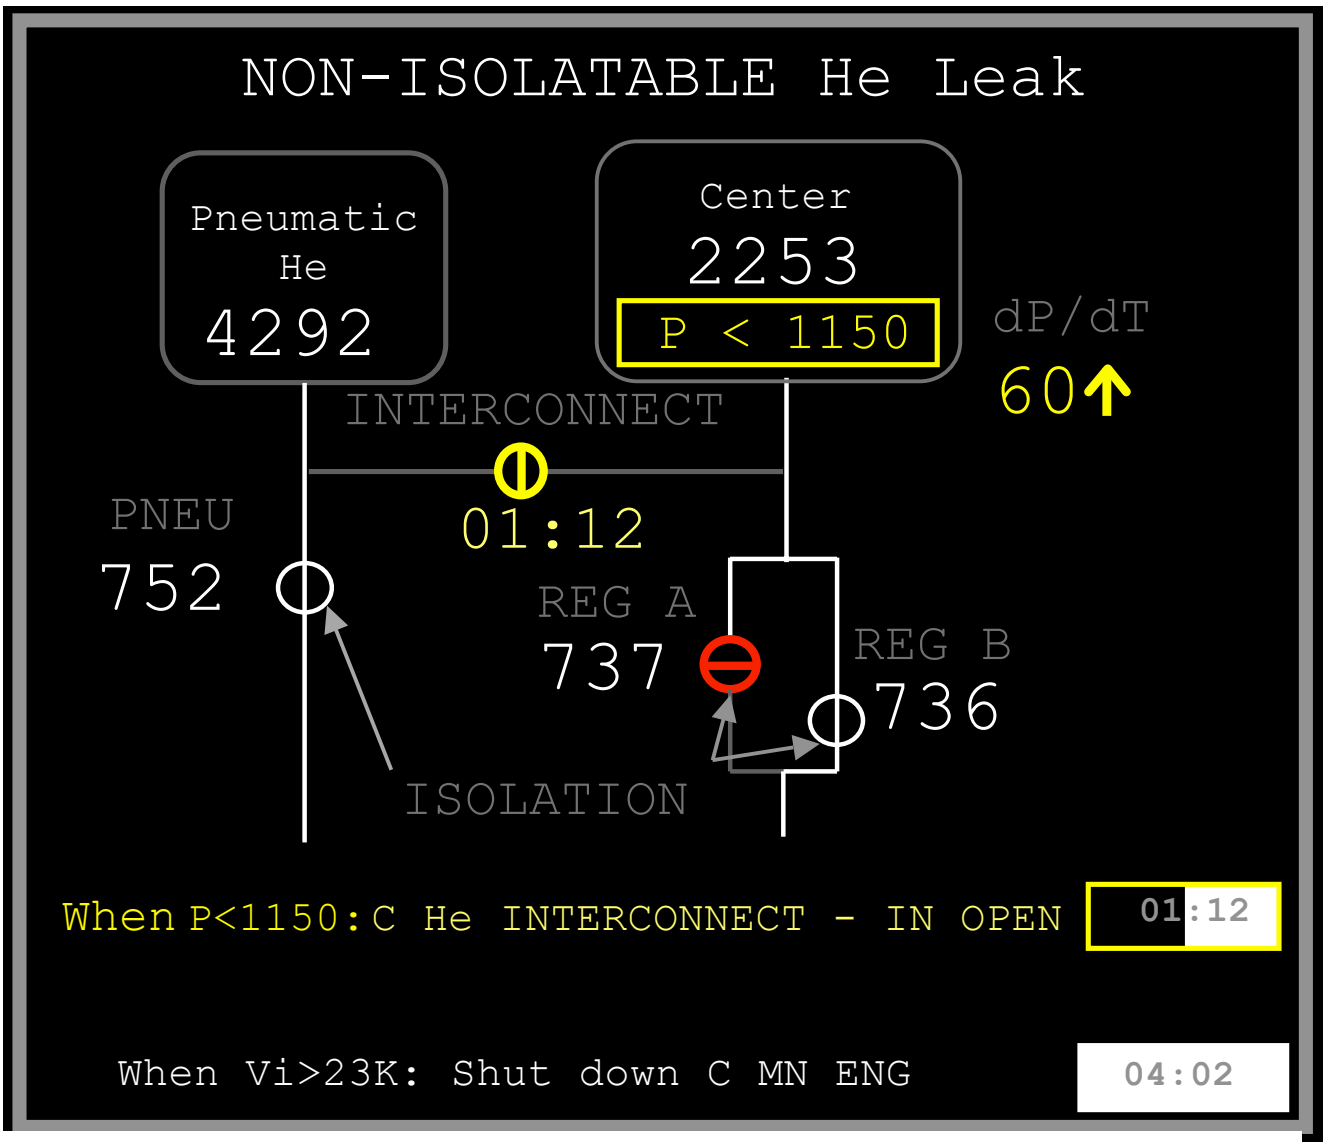

Figure 3: Candidate Fault Management Display in "Countdown" Mode.

"countdown" timer is provided just to the right of the written instruction showing 1 min 12 sec until this procedure should be performed.

Meanwhile, the upper half of the display contains a schematic of the affected system and its components. Since the interconnect procedure is going to be performed first, the schematic shows a "blow-up" of the affected components of the helium supply system. The procedural information is embedded in the schematic via size-and-color-coding: the interconnect valve is enlarged and yellow, and a yellow countdown timer appears below the interconnect valve symbol. The schematic always depicts the current operational configuration, so the interconnect valve is shown in the closed position with no flow through the interconnect manifold. The color-coding signals that the upcoming procedure will command a change from the current (closed) position to the alternative (open) position.

When the condition needed to carry out this activity is satisfied, the fault management display changes to "command" mode (Figure 4). The countdown symbols disappear, and the Interconnect valve symbol turn magenta. In parallel, the written procedure also turns magenta, and a virtual magenta "Accept" button appears to the right of the procedure. Again, referring to the schematic, magenta signals a recommended change from current (closed) to alternative (open) position. After review, the crewmember signals his or her agreement to proceed with the "Open Interconnect" procedure by touching the "Accept Button" to the right of the written procedure. Once the automation has performed the action (not shown in the Figure), the "He Interconnect" procedure shifts down, and turns grey,

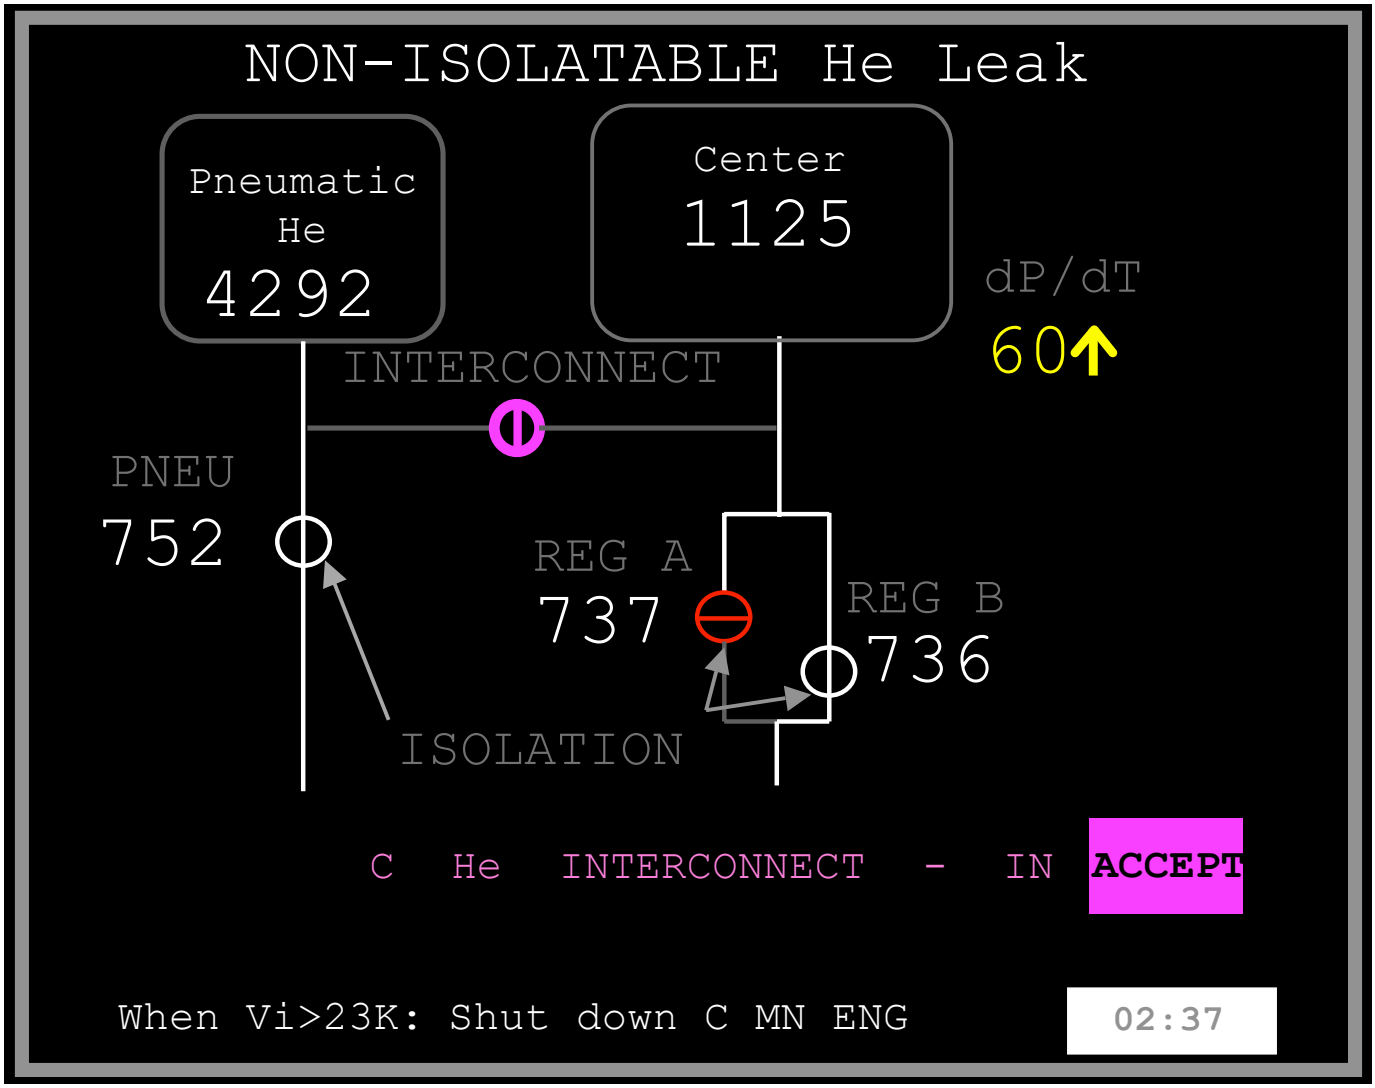

Figure 4: Candidate Fault Management Display in "Command" Mode

and the engine shutdown line moves to the top of the stack. The schematic converts to a main engine synoptic showing the crucial valves that must be closed in order to achieve Engine Shutdown.

The "permissions" mode of human-automation interaction, and this user interface prototype, is designed to keep the crewmember in close synchronization with the automation and closely track the automated actions. On the upper (schematic) section, the slight mental exercise required to translate current valve position into the commanded position should further enhance situation awareness of the nature of the malfunction and the system configuration that will result from the recommended action.

5. Preserving Backup Capability with Head-Down Displays.

This interface concept still contains a significant human factors drawback. Results of a recent study of space shuttle fault management behavior by relatively novice operators[14] has provided strong evidence that knowing the locations of the switch panels and their embedded switches plays an important role in enabling crewmembers to work malfunctions quickly and accurately. By requiring crewmembers to locate and throw the switches themselves (current operations), cockpit training in ground-based simulators continuously preserves and reinforces this knowledge. By contrast, the proposed system does not require the crewmember to physically locate the switches, potentially degrading their ability to do so in a CEV emergency.

Could we redesign the user interface to gain the benefits of automatic procedure execution, but preserve the crew's backup capability? One possibility might be to overlay each of the cockpit switch panels with a "head-down display", a (removable or retractable) layer of glass on which arbitrary forms of information and symbology could be superimposed. Since the health manger performs all the actual procedures, the actual hardware switches underneath the glass panels remain in their "Computer-Controlled" position (usually, the middle position of a

standard up-down switch configuration). The head-down displays would include "virtual" switches that show actual switch positions. When a procedure calls for a switch throw, the change could be signaled on the virtual switch panel via color-coding that makes the affected switch "pop out" at the observer from the virtual display. And, by having the "accept action" button appear right beside the commanded switch, the crewmember would be forced to orient to the location of the switch in the cockpit, constantly reinforcing his or her spatial representation of the position of the switches associated with particular systems and particular operations. The crewmember would therefore retain the knowledge to take over and throw the hardware switches manually, if required.

The idea of overlaying the cockpit switch panels with glass displays has other potential benefits. Since all the hardware switches would be in (and remain in) the "Computer-Controlled" position, the physical switch panels would only provide information concerning actual (current) valve positions (e.g., open, closed, or in transition) if talkback indicators are present. The necessity to include a talkback indicator for every switch greatly increases the real estate requirements of the switch panels. If the switch panels were overlaid by head-down displays, actual switch positions and talkback indicators could all be depicted right on top of the switch panels, greatly reducing these requirements.

Finally, by distributing procedural information across two separate cockpit locations (on top of the switch panels and on the dedicated fault management display), we have a chance to build in some crosschecking capability similar to the crosschecking between the Primary and Backup flight software systems. We have already noted that reasoning systems come in many different forms. Suppose each systems-level health manager incorporated two distinct reasoning systems that worked redundantly and independently on root-cause determinations and fault diagnosis. One system could be generating root-cause fault determinations, and providing the appropriate procedures, on the fault management display. The other system could be doing parallel computations, also leading to the specification of the appropriate procedures, but depicting them on the head-down displays. Crosschecking the recommended procedures across the two display formats would enable the crew to assess the veracity of the information being generated by the fault management systems, and increase crew trust in the recommended actions. The cross checking requirement would also continuously reinforce crewmembers' spatial knowledge of actual switch locations, again helping maintain the spatial knowledge needed to execute procedures manually in the event of hardware/software failures.

III. Enhancing crew mission management capabilities: Multi-Modal Interfaces

Thus far, our discussion has focused on how to exploit the ever-growing information processing resources of portable computing devices to enhance onboard mission management capability. We did not step outside the conventional approach to human-machine interaction whereby virtually all information sources are visual, and virtually all physical interfaces between the crewmembers and the vehicle (e.g., keyboards, button presses, switch throws, etc.) are operated with the hands and fingers. However, user interface technologies have now progressed to the point where information can also be presented in the form of auditory and haptics "displays", and natural language understanding systems make it possible to operate machines by voice command as well as by hand. To the extent that people can processes information from multiple modalities in parallel, and formulate and execute manual actions in parallel with vocal responses, incorporating these nontraditional user interfaces in ETS vehicle cockpits could significantly enhance the crew's mission management capabilities, particularly during dynamic flight phases when processing demands are highest.

Studies of human information processing capabilities have established the existence of a "preattentive" processing mode that extracts information from all modalities in parallel and automatically directs focal attention to stimuli that are salient to the operator's current task set[15]. Woods[16] provides a relevant example of such a stimulus, which he calls a "preattentive reference", from the control room of a nuclear power plant. In this environment, operators function as process controllers, monitoring and troubleshooting complex engineering systems for signs of off-nominal functioning. As in today's spacecraft cockpits, systems information is provided almost exclusively on visual displays, so operators are continuously scanning these displays and instrument readings for off-nominal readings. In this particular control room, incidental auditory cues in the form of discrete "clicks" were present, and the click rate carried useful information about fuel rod status. An unusual "click rate", indicating a disturbance in nominal status, automatically captured operators' attention even while they were concentrating on their visual displays.

In a next-generation spacecraft cockpit, preattentive references could be deliberatively built in to a multi-modal display system, enhancing a crewmember's ability to monitor vehicle systems and flight parameters compared to visual displays alone. A shuttle-based example for ascent would be to generate a set of spatially distinct tones whose perceived 3-D (perceived) location corresponds to the standard labels for the three main engines (Left, Center, and Right) and whose amplitude (loudness) corresponds to main engine thrust level. Amplitude-based thrust coding

represents a 'sonification' approach to synthesizing sound cues in which critical operational parameters are mapped to acoustic features easily discernible by the listener[17]. All three engines maintain the same thrust profile from liftoff to engine cut-off, approximately 8.5 min into flight. If the amplitude of one tone changed suddenly relative to the others, the operator would be quickly alerted to the presence of an abnormal thrust level and with the spatial coding, to which engine was affected, all without having to look directly at the main engine summary display.

Beyond these alerting qualities, it is not clear to what extent multi-modal interfaces would enhance a crewmember's capability to work a complex onboard operation, such as fault detection, isolation, and recovery. One can imagine a multi-modal system in which a crewmember would work one malfunction via the standard "visual input – manual output" channel, and a second, unrelated, malfunction through a natural language interface. The health management system for the second malfunction would provide relevant information, including root-cause identification and procedures, via spoken language, and the crewmember would give permission to execute the procedures via voice commands. There are no obvious peripheral bottlenecks that would prevent this auditory-vocal "channel" from functioning fully in parallel with the traditional "visual-manual" channel. In principle, then, the crewmember could work both malfunctions as quickly and efficiently as working just one.

Unfortunately, while it is well established that humans have considerable parallel processing capabilities at the perceptual and motor (i.e., peripheral) stages, there is much less agreement concerning how much parallel processing capability exists for more central operations, such as the decision-making involved in working a malfunction. Some cognitive processing models assume that even complex forms of multi-tasking, such as malfunction handling, should be possible[18,19]. Other models[20,21,22] suggest that our complex multi-tasking capabilities are sharply limited by a central processing "bottleneck" that ought to limit decision-making and other central cognitive operations to only one task at any one time. So little is known about human processing capabilities in complex multi-tasking environments that Woods has called the issue "the least explored frontier in cognitive science and human-machine cooperation"[23]. Certainly, without a better understanding of how multi-modal information displays and controls affect human parallel processing capabilities, we will not be able to determine the appropriate role for multi-modal interfaces in ETS cockpits.

IV. Human-Automation Issues for Long-Duration Missions.

G. Adaptive Cockpits

Earlier in the paper, we noted that automating mission operations increases the requirements on crewmembers to act as backups in case of hardware or software failure. We also made passing reference to the reverse case, where automation would act as a backup for a "human failure", taking over and accomplishing functions that crewmembers normally perform in the event they are incapacitated. In our concept for mixed initiative fault management, for example, this requirement would involve adjusting the automation level "upward" to the point where the automated health manager could execute procedures without crew permission.

This kind of adaptive capability may not be a priority for the early (LEO) phase of ETS operations, where mission durations will be short, crewmembers are likely to be functioning at a high level, and plenty of real-time ground assistance is available. As the missions increase in duration, however, requirements for adaptive capabilities will grow. On these missions, crewmembers will experience long-term exposure to various space-based environmental stressors, such as circadian disruptions (fatigue), confinement, microgravity, and possibly elevated doses of radiation. These stressors have considerable potential to impact crewmembers' operational capabilities. During the quiescent (cruise) phase, any performance decrements will not usually form much of a mission risk, as there are few situations where information processing and decision-making requirements are sufficiently high to stress human capabilities. However, these periods are always followed by a highly dynamic flight phase where, for a short time, crewmembers are called upon to manage and participate in activities that *do* make strong demands on their information-processing resources. In a relatively recent interview, Neil Armstrong identified piloting and landing the lunar excursion module as by far the highest workload phase of his Apollo mission (despite the fact that he was receiving real-time assistance from the ground). On missions to Mars, the demands placed on the crew will be even greater than Apollo, since the crew will have to take full responsibility for vehicle operations.

Unfortunately, the abrupt transition to these high-pressure operational environments will be accompanied by a transition to a new gravitational environment. Descent and landing to the Martian surface will be accompanied by a transition from several months of 0-G to a variable and fractional G. Liftoff from the surface, followed by Rendezvous and Docking in Mars orbit, will be accompanied by a transition from several weeks in a steady .38 G environment to another (and temporally variable) fractional-G environment; and last but not least, Entry, Descent, and Earth Landing will be accompanied by a transition from several further months of 0-G environment to high-G.

Deconditioning and adaptation issues arising from abrupt G transitions are already known to impact human performance; this knowledge lies behind the operational decision to confine shuttle missions to 14 days or less.

To ensure optimal crew-system performance during the safety-critical flight phases, a strong argument can be made that long-duration missions have unique requirements for a cockpit or crewstation environment that can accurately assess the current state (i.e., level of fatigue, deconditioning, cognitive capabilities) of individual crewmembers, and then *adapt* to that state, in real time. Adaptive requirements run the gamut from what information should be displayed to what crewmember, to the modality in which the information is presented, to the choice of functional allocation between crewmembers (e.g., task assignments, task schedules), to the functional allocation between crewmembers and onboard automation (i.e., adaptive/adjustable level of automation).

Several technologies would have to be further developed and integrated in order to develop a reliable and effective adaptive system. Real-time inference tools will have to be developed that can reliably determine a crewmember's cognitive state and cognitive functioning from electrophysiological measures and behavioral measures, such as cockpit eye scan patterns[2]. Similar to a advanced health monitoring system for a machine, the inference tool would have to be able to monitor these human performance data feeds in real time, and accurately classify (diagnose) any anomalous behavior as arising from a particular stressor (fatigue, deconditioning, adaptation, etc.) or set of stressors. Once this determination (diagnosis) was made, an appropriate mitigation strategy could be selected, involving environment adaptation along the various dimensions identified in the previous paragraph.

Researchers at NASA Ames Research Center are starting to develop human performance models of crewmember behavior. In the future, these models could be customized to encode knowledge of crew-specific behaviors and provide crew-specific determinations of current capabilities and appropriate mitigation strategies.

H. Human-Automation Interaction During Cruise

The need for a software "agent" that encodes crewmember-specific knowledge, and then uses that knowledge to customize crew-vehicle interactions, is not confined to dynamic flight phases. During the long-duration cruise phases, crewmembers will typically be spatially distributed throughout the vehicle. At any one time, they will be engaged in any one of a number of heterogeneous activities. If off duty, they might be sleeping, eating, or attending to personal hygiene. If on duty, they might be operating and/or maintaining the various complex on-board systems and subsystems that work in continuous mode, such as the electrical power system, environmental control, food processing, air/water recycling, and solid waste processing. Many of these activities are asynchronous, in the sense that an activity or group of activities may be started by one crewmember during one shift and then finished by another crewmember on the following shift. Such "loosely-coordinated" forms of group interaction carry unique requirements for planning, scheduling, and coordinating group activities and for enabling efficient interactions among the various members of the group. As the moment-by-moment operations of the onboard systems become increasingly reliant on ("smart") systems controllers (such as the health manager depicted in Figure 2), the activity coordination has to confront the problem of coordinating team activities when some of the team members are software agents.

Researchers at, or affiliated with, NASA's Johnson Space Center have developed an architecture for sharing management and maintenance duties between controllers and groups of physically distributed humans[24]. The cornerstone of their Distributed Collaboration and Interaction (DCI) architecture is the Attentive Remote Interaction and Execution Liaison (ARIEL) agent, of which there is one for each crewmember. Each agent models and keeps track of their crewmember's location, activity, group role, and availability, to help the crewmember in communicating and coordinating his/her activities within the group. The agent performs a variety of roles to support and enable effective collaboration between individual crewmembers and onboard automation, on an as-needed basis. Thus, the final human-automation challenge we will identify is the need to integrate ARIEL agents with the human performance model-based agents we propose to support adaptive cockpits during dynamic flight phases. That way, communications and other protocols for human-automation interaction onboard ETS vehicles can be standardized across all mission and flight phases.

Acknowledgments

This work was supported by NASA's Space Human Factors Engineering Program (RTOP # 131-20-30). We thank members of the Intelligent Spacecraft Interface Systems (ISIS) Lab at NASA Ames Research Center for useful discussion and assistance with display design.

References

[1]Huemer, V. A., et al., "Characterizing Scan Patterns in a Spacecraft Cockpit Simulator: Expert versus Novice Performance", *HFES 49th Annual Conference* (in press).

[2]Hayashi, M, Beutter, B., & McCann, R. S. Hidden Markov Model Analysis for Space Shuttle Crewmembers' Scanning Behavior. *IEEE 2005 International Conference on Systems, Man, and Cybernetics* (in press).

[3]Rogers, W. H., Schutte, P. C., & Latorella, K. A. Fault Management in Aviation Systems. In R. Parasuraman & M. Moulsua (Eds*.), Automation and Human Performance: Theory and Application*., Erlbaum, New Jersey, pp. 281-317.

[4]Williams, B. C., & Nayak, P. P. "Immobile Robots: AI in the New Millenium. *AI Magazine*, Fall 1996.

[5]Malin, et al., "Making Intelligent Systems Team Players: Case Studies and Design Issues, Vol. 1: Human-Computer Interaction Design," NASA TM #104738, 1991.

[6]Billings, C. E., *Aviation Automation: The Search for a Human-Centered Approach*., Erlbaum, Hillsdale, New Jersey, 1997.

[7]Endsley, M. R., & Kiris, E. O., "The out-of-the-loop performance problem and level of control in automation*", Human Factors*, Vol. 37, 1995, pp. 381-394.

[8]Malin, J. T., Schreckenghost, D. L., & Rhoads, R. W., "Making Intelligent Systems Team Players: Additional Case Studies," NASA TM #104786, 1993.

[9]Scandura, P. A., & Garcia-Galan, C. A., "A Unified System to Provide Crew Alerting, Electronic Checklists and Maintenance Using IVHM," *IEEE DASC Conference*, CP312, 2004.

[10]McCann, R. S., & McCandless, J., "Human-Machine Teaming for Dynamic Fault Management in Next-Generation Space Vehicles", *JANNAF 3rd Modeling and Simulation Subcommittee CP*, 2003.

[11]Huff, E. M., Tumer, I. Y., & Mosher, M., "An Experimental Comparison of Transmission Vibration Responses from OH-58 and AH-1 Helicopters", AHS 57th Annual Forum, 2001.

[12]McCandless, J. W., et al., "Evaluation of the Space Shuttle Cockpit Avionics Upgrade (CAU) Displays", *HFES 49th Annual Conference* (in press).

[13]Malin, J., et al., "Multi-agent Diagnosis and Control of an Air Revitalization System for Life Support in Space", *IEEE Aerospace CP*, 2000.

[14]Huemer, V. A., Matessa, M. P., & McCann, R. S., "Fault Management during Dynamic Space Flight: Effects of Cockpit Display Format and Workload", *IEEE 2005 International Conference on Systems, Man, and Cybernetics* (in press).

[15]Folk, C. L., Remington, R. W., & Johnston, J. C., "Involuntary covert orienting is contingent on attentional control settings", *Journal of Experimental Psychology: Human Perception and Performance*, Vol. 18, 1992, pp. 1030-1044.

[16]Woods, D., "The alarm problem and directed attention in dynamic fault management", *Ergonomics*, Vol. 18, 1995, pp. 2371-2393.

[17]Kramer, G. (ed.), *Auditory Display, Sonification, Audification, and Auditory Interfaces*, Addison-Wesley, Reading, MA, 1994.

[18]Wickens, C. D., Sandry, D. L., & Vidulich, M., "Compatibility and Resource Competition Between Modalities of Input, Central Processing, and Output," *Human Factors*, Vol. 25, 1983, pp. 227-248.

[19]Meyer, D. E., & Kieras, D. E., "A Computational Theory of Executive Cognitive Processes and Multiple-Task Performance: Part 1: Basic Mechanisms," *Psychological Review*, Vol. 104, 1997, pp. 749-791.

[20]Pashler, H., "Processing Stages in Overlapping Tasks: Evidence for a Central Bottleneck," *Journal of Experimental Psychology: Human Perception and Performance*, Vol. 10, 1984, pp. 358-377.

[21]McCann, R. S., & Johnston, J. C., "Locus of the Single-Channel Bottleneck in Dual-Task Interference*," Journal of Experimental Psychology: Human Perception and Performance*, Vol. 18, 1992, pp. 471-484.

[22]Lien, M.C., McCann, R. S., Ruthruff, E., & Proctor, R. W., "Dual-Task Performance with Ideomotor-Compatible Tasks: Is the Central Processing Bottleneck Intact, Bypassed, or Shifted in Locus?*," Journal of Experimental Psychology: Human Perception and Performance*, Vol. 31, 2005, pp. 123-144.

[23]Woods, D. D., Johannesen, L. J., Cook, R. L,, & Sarter, N. B., "Behind human error: Cognitive Systems, Computers and Hindsight", Crew Systems Ergonomic Information and Analysis Center (CSERIAC) State of the Art Report, Dayton, OH, 1994.

[24]Schreckenghost, D., et al., "Supporting Group Interaction Among Humans and Autonomous Agents," NASA Johnson Space Center TRACLabs Report, Houston, TX, 2004.